



(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention
of the grant of the patent:
31.10.2001 Bulletin 2001/44

(51) Int Cl.7: **H04Q 11/04, H04L 12/56**

(21) Application number: **00301164.0**

(22) Date of filing: **15.02.2000**

(54) **Computationally-efficient traffic shaper**

Vom Rechenaufwand her effizientes Verkehrsformer

Appareil de mise en forme du trafic économisant le volume calculatoire

(84) Designated Contracting States:
DE FR GB

(30) Priority: **23.02.1999 US 256015**

(43) Date of publication of application:
20.09.2000 Bulletin 2000/38

(73) Proprietor: **LUCENT TECHNOLOGIES INC.**
Murray Hill, New Jersey 07974-0636 (US)

(72) Inventor: **Petty, Norman W.**
Boulder, Colorado 80302 (US)

(74) Representative: **Williams, David John**
Page White & Farrer,
54 Doughty Street
London WC1N 2LS (GB)

(56) References cited:

- **"WEIGHTED QUEUEING ALGORITHM FOR EFFICIENT ASYNCHRONOUS TRANSFER MODE TRAFFIC SHAPING" IBM TECHNICAL DISCLOSURE BULLETIN, US, IBM CORP. NEW YORK, vol. 39, no. 4, 1 April 1996 (1996-04-01), pages 161-163, XP000587459 ISSN: 0018-8689**
- **KALMANEK C R ET AL: "A HIGH-PERFORMANCE QUEUEING ENGINE FOR ATM NETWORKS" PROCEEDINGS OF THE INTERNATIONAL SWITCHING SYMPOSIUM, JP, TOKYO, IEICE, vol. SYMP. 14, 1992, pages 354-358, XP000337673**

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

[0001] This invention relates to traffic shapers for, and to methods of shaping traffic of, a plurality of streams of traffic.

[0002] Today's business communications environment consists of two separate network infrastructures: a voice network (such as a private branch exchange (PBX)) characterized by real-time, high-reliability, constant bit-rate (CBR) connections; and a data network (such as a packet network) characterized by high-bandwidth variable bit-rate (VBR) connections. Business needs for simplified maintenance, management, and access to information on diverse networks are forcing the convergence of these networks along with a new class of real-time multimedia networks. Asynchronous transfer mode (ATM) provides a single infrastructure that cost-effectively and flexibly handles both switching and transmission for the traffic types mentioned above (voice, video, and data) for both local-area networks and wide-area networks. The evolving network convergence requires the adaptation of the legacy PBX voice traffic to ATM. Voice telephony over ATM (VTOA) specifications allow adaptation of compressed or uncompressed voice pulse-code modulated (PCM) data streams into streams (virtual circuits) of CBR cells.

[0003] An ATM cell, regardless of the traffic it carries, is a packet 53 octets long: 48 octets of payload attached to a 5-octet header. The header contains addressing and management information used to direct the cell from source to destination and to ensure that the negotiated aspects of the traffic-flow through the ATM network are met. CBR traffic is assembled into cell payloads using ATM Adaptation Layer 1 (AAL1). The AAL1 cell constructor layer uses the first octet of the payload for its header and the remaining 47 octets to carry CBR information. ATM cell construction is then completed by attaching the ATM header to the payload.

[0004] An individual ATM cell carries traffic of a single narrowband or wideband channel (a virtual circuit). Illustratively, a narrowband channel is represented by a single time slot of each successive frame of a TDM bus, while a wideband channel is represented by a plurality of time slots of each successive frame of a TDM bus. If a transmitter handles a plurality of virtual circuits, different ones of the ATM cells that it transmits carry traffic of different virtual circuits. Likewise, if a receiver handles a plurality of virtual circuits, different ones of the ATM cells that it receives carry traffic of different virtual circuits. A transmitter may transmit the traffic of different virtual circuits to different receivers. Likewise, a receiver may receive the traffic of different virtual circuits from different transmitters.

[0005] C.R. Kalmanek et al., "A High-Performance Queuing Engine for ATM Networks", International Switching Symposium 1992, (ISS '92, 25-30 October 1992, Yokohama, Japan) Proceedings Vol. 1, pp. 354-358, XP000587459, discloses a queuing engine for

an ATM switch, which has two queues dedicated to each ATM virtual circuit. An input queue receives ATM cells from a trunk and stores them for transmission on a back-plane to an address translation module. An output queue receives ATM cells from the address translation module over the back plane and stores them for transmission on a trunk. A service decoder controls the serving order of queues by controlling which prioritized service list they are put on. A queue server serves the queues on each list on a round-robin basis.

[0006] The traffic of each virtual circuit is processed by its own instance of the AAL1 cell constructor layer; the number of active instances varies as virtual circuits are added and removed. Assuming a channel carrying uncompressed pulse-code modulated (PCM) traffic having an 8KHz sampling rate, each instance of the cell constructor that is processing a narrowband virtual circuit is supplied one PCM octet every 125us frame period. It therefore requires 5.875ms to fill a cell (47 octets * 125us/octet). Fully constructed (mature) cells from all instances of the cell constructor are then transmitted.

[0007] Unlike the constant cell rate from narrowband virtual circuit, a wideband virtual circuit generates a constant average cell rate but a variable instantaneous cell rate. For example, a 32-narrowband-channel-wide wideband virtual circuit is supplied 32 octets during each frame period. The processing sequence effected by its cell constructor instance is as follows: during a first frame period, fill the first 32 octets of a first empty ATM cell; during a second frame period, fill the remaining 15 octets of the first ATM cell, transmit the first ATM cell, and fill the first 17 octets of a second empty ATM cell; during a third frame period, fill the remaining 30 octets of the second cell, transmit the second cell, and fill the first 2 octets of a third empty cell; etc. This sequence shows an instantaneous cell rate of one cell every 250us (two frame periods) sometimes and one cell every 125us (one frame period) at other times. Variable bit-rate virtual circuits, whether narrowband or wideband, provide even greater diversity in the instantaneous cell rate.

[0008] ATM switches police cell traffic to ensure that it does not exceed the cell-traffic capacity of the switches. The average cell rate is policed on a per-switch basis, but the peak (instantaneous) cell rate is policed on a per-virtual-circuit basis. Receiving ATM switches delete received cells of each virtual circuit that exceed that virtual circuit's negotiated peak cell rate. And transmitting ATM switches meter out cells from each virtual circuit to ensure that they do not exceed that virtual circuit's negotiated peak cell rate.

[0009] The ATM standards specify a Generic Cell Rate Algorithm (GCRA) for use by both hardware and software-implemented devices, called traffic shapers, for metering out cells. The GCRA requires the traffic shapers to periodically process the output queue of each virtual circuit. The period may be different for each virtual circuit, and is a function of the negotiated peak

cell rate for that virtual circuit. In the case of PCM traffic having an 8KHz sampling rate, the processing period is the 125us frame period. The queues of all virtual circuits are processed during each period. If the queue of a virtual circuit is empty, the traffic shaper does not transmit a cell from that queue. If a cell is available in the queue of a virtual circuit, the traffic shaper transmits it if it meets the GCRA criterion, i.e., if the period between the last transmitted cell from this queue and now is greater than or equal to the minimum administered cell-rate period. If more than one cell is available in the queue of a virtual circuit, the traffic shaper transmits one cell from the queue if it meets the GCRA criterion, and the next cell must wait at least until the subsequent processing period to be transmitted. A spacing of at least 125us between transmitted ATM cells from any virtual circuit is thus ensured, thereby guaranteeing that the virtual circuit does not exceed its negotiated peak cell rate.

[0010] A problem with using the GCRA for multiple virtual circuits is that the traffic shapers must test the output queue of each virtual circuit during each processing period to determine if it has a cell to transmit. This is a particularly onerous task for GCRA implemented in software that supports large numbers (e.g., hundreds) of virtual circuits. Statistically, however, only one in 47 GCRA queue checks will result in the transmission of a cell during each processing period.

[0011] According to one aspect of this invention there is provided a traffic shaper as claimed in claim 1.

[0012] According to another aspect of this invention there is provided a method of shaping traffic as claimed in claim 6.

[0013] Generally, traffic shaping for a plurality of streams of traffic (e.g. ATM virtual circuits or other packetized channels) utilizes a sequence of a first plurality of queues that are shared by the plurality of traffic streams, an enqueueing arrangement, and a dequeuing arrangement. Each of the traffic streams has at least one of the queues designated as its serving queue. Dequeueing is effected by cyclically transmitting contents of the sequence of the queues, i.e., by transmitting the contents of a different sequential one of the queues during each sequential period of the first plurality of such periods. During each period, a queue having a fixed offset in the sequence from the queue whose contents are being transmitted forms a current queue for that period. Preferably, the offset is zero, i.e., the current queue is the one whose contents are being transmitted. Enqueueing is effected as follows. In response to receipt of traffic (e.g., an ATM cell or another packet) from any said traffic stream during a period before a serving queue of that traffic stream becomes the current queue for that period, the received traffic is stored in that serving queue of that traffic stream. In response to receipt of traffic from any said traffic stream during a period when a serving queue of that traffic stream is the current queue for that period, the received traffic is stored in the current queue for that period. In response to receipt of traffic from any said traf-

fic stream during a period after a serving queue of that traffic stream ceased being the current queue for that period, the received traffic is stored in the current queue for that period, and the designation of the at least one serving queue of that traffic stream is changed to designate the current queue of that period as a serving queue of that traffic stream. If a traffic stream has a plurality of serving queues, they are evenly spaced apart in the sequence of queues. The designation of the serving queues is then changed by shifting the designation from each of the serving queues a same distance in the sequence, to other queues that include the current queue. The shifting preferably involves shifting the designation a distance in the sequence from the serving queue of that traffic stream that last ceased being the current queue to the current queue, and shifting the designation from each of the other serving queues of that traffic stream a same distance to another queue.

[0014] This arrangement effects traffic shaping for both narrowband and wideband constant-bit-rate channels. For variable-bit-rate channels, an additional overflow queue is used, for temporarily storing any traffic received when enqueueing that traffic in the first plurality of queues would result in storing of that traffic in a queue that presently already stores received traffic of that traffic stream. When dequeuing has made room for storing the traffic that is presently stored in the overflow queue available in the first plurality of queues, enqueueing of that traffic in the first plurality of queues is then effected.

[0015] The above-characterized arrangement is very computationally-efficient. For example, while the GCRA of the prior art requires that as many queues as there are traffic streams be checked for an item of traffic to transmit during each dequeuing period, the subject arrangement requires that only one queue be checked and its contents be transmitted during each dequeuing period. Moreover, the subject arrangement often reduces the number of queues that are needed to effect traffic shaping. For example, in the case of an ATM system that transports only narrowband constant-bit-rate virtual circuits, only 47 queues are needed, regardless of how many of the virtual circuits there are. Although it is described below as providing traffic shaping for a plurality of ATM virtual circuits, it can be used for any packetized channels, including frame relay and voice over the Internet, and in fact is not precluded from serving non-packetized channels as well. And although it can be implemented in hardware, the subject arrangement is particularly suited for implementation in software, including firmware.

[0016] These and other features and advantages of the present invention will become more apparent from the following description of an illustrative embodiment of the invention considered together with the drawing.

Brief Description of the Drawing

[0017]

FIG. 1 is a block diagram of an ATM cell constructor that includes an illustrative embodiment of the invention;
 FIG. 2 is a diagram of a first embodiment of a traffic shaper of the ATM cell constructor of FIG. 1;
 FIG. 3 is a functional flow diagram of a global initialization function of a state machine of the traffic shaper of FIG. 2;
 FIG. 4 is a functional flow diagram of an instance initialization function of the state machine of the traffic shaper of FIG. 2;
 FIG. 5 is a functional flow diagram of an enqueueing function of the state machine of the traffic shaper of FIG. 2;
 FIG. 6 is a functional flow diagram of a dequeueing function of the state machine of the traffic shaper of FIG. 2;
 FIG. 7 is a diagram of a second embodiment of the traffic shaper of the ATM cell constructor of FIG. 1;
 FIG. 8 is a functional flow diagram of an instance initialization function of a state machine of the traffic shaper of FIG. 7;
 FIG. 9 is a functional flow diagram of an enqueueing function of the state machine of the traffic shaper of FIG. 7;
 FIG. 10 is a functional flow diagram of a dequeueing function of the state machine of the traffic shaper of FIG. 7;
 FIG. 11 is a diagram of a third embodiment of the traffic shaper of the ATM cell constructor of FIG. 1;
 FIG. 12 is a functional flow diagram of an enqueueing function of the state machine of the traffic shaper of FIG. 11; and
 FIG. 13 is a functional flow diagram of a dequeueing function of the state machine of the traffic shaper of FIG. 11.

Detailed Description

[0018] FIG. 1 shows an ATM cell constructor 100, also known as an ATM cell assembler, such as may be used in an interface port circuit of a PBX or in any other ATM interface apparatus to construct ATM cells from a stream of traffic, such as voice and/or video traffic.

[0019] Cell constructor 100 and each of its components may be individually implemented either in hardware or in software/firmware, either alone or in an integrated circuit with other devices. The software or firmware may be stored in any desired storage device readable by a computer—for example, a read-only memory (ROM) device readable by an interface port circuit processor. Multiple virtual circuits (also referred to herein as channels, calls, or communications) of CBR traffic are received by ATM cell constructor 100 over a

communications medium 102, and follow a data path 149 through ATM cell constructor 100 where successive segments of the traffic are formed into packets (ATM cells). If the switching system employing ATM cell constructor 100 is the Definity® PBX of Lucent Technologies Inc., medium 102 is a time-division multiplexed (TDM) bus that carries up to 242 individual channels of traffic in 242 individual time slots of repeating frames. Each frame carries one (narrowband) or more (wideband) time slots of each channel's traffic stream. Each time slot carries one byte (octet) of traffic.

[0020] The bytes of traffic of individual channels are assembled into ATM cells in data path 149. It takes on the order of a TDM bus frame-interval to process an individual time slot of traffic through data path 149; of course, up to a frame's worth of time slots may be processed in parallel. A TDM bus frame-interval is therefore taken as a cell construction period. It is a predetermined time interval during which each virtual circuit can mature an ATM cell for transmission. It can take up to 47 frames to construct a cell, however. An ATM processor 118 sequentially receives mature ATM cells and transmits them on an ATM communications medium 120 towards their destinations. ATM processor 118 comprises a conventional ATM layer 122 and a conventional physical layer 123 interfaced by a traffic shaper 121. Traffic shaper 121 ensures that the peak instantaneous cell transmission rate of each virtual circuit does not exceed its negotiated peak cell rate. Traffic shaper 121 comprises a sequence 130 of shaping queues 131, a traffic shaper dequeue state machine 141, and a plurality of instances of traffic shaper enqueue state machine 140, one instance for each virtual circuit.

[0021] A first illustrative traffic shaper 121 embodying the invention is shown in FIG. 2. This embodiment is tailored to serve narrowband CBR virtual circuits. It replaces the queue-per-virtual-circuit of the prior art with a sequence 130 of 47 shaping queues, designated in FIG. 2 as Q0-Q46, one per AAL1 cell construction state (i.e., one per ATM cell payload octet), which are shared by all of the virtual circuits.

[0022] During each queue processing period (which is a time interval at most equal to an inverse of a peak instantaneous cell rate of a fastest one of the virtual circuits, which is the frame period of 125us in this instance), one of the queues 131 - the current queue - is emptied by dequeue state machine 141 and its contents are sent to physical layer 123. Alternatively, instead of the queue that is presently being emptied, a queue that has a fixed offset of one or more queues 131 in the sequence 130 may be referred to as the current queue during each processing period, for purposes of operation of enqueue state machine 140. During successive sequential queue processing periods, successive sequential ones of the queues 131 are designated as the current queue. Hence, it takes 5.875ms (47.125us) to empty all queues 130.

[0023] Each instance of enqueue state machine 140

identifies the one of the queues 131 that is designated as serving its corresponding virtual circuit (VC). In the example of FIG. 2, Q0 is indicated as serving VC 4, 12, and 13; Q1 is indicated as serving VC 7, 9, and 10; Q2 is indicated as serving VC 1, 5, 11, and 14; Q3 is indicated as serving VC 6 and 16; Q4 is indicated as not serving any VC; Q5 is indicated as serving VC 15; etc; and Q 46 is indicated as serving VC 3 and 14. When ATM layer 121 makes an ATM cell of a virtual circuit available, that virtual circuit's corresponding instance of enqueue state machine 140 enqueues the cell in one of the queues 131, selected as follows. If the cell is on-time, meaning that it becomes available while the serving queue of this VC is the current queue, the state machine instance enqueues the cell in the current queue. For example, if a cell of VC 5 becomes available while Q2, VC 5's serving queue, is the current queue, the cell is enqueued in Q2, as shown in FIG. 2. If the cell is early, meaning that it becomes available before the serving queue of this virtual circuit becomes the current queue, the state machine instance enqueues the cell in the serving queue of this virtual circuit. For example, if a cell of VC 14 becomes available while Q1 is the current queue, the cell is enqueued in Q2, VC 14's serving queue; similarly, if a cell of VC 6 becomes available while Q2 is the current queue, the cell is enqueued in Q3, VC 6's serving queue, as shown in FIG. 2. If the cell is late, meaning that it becomes available after the serving queue of this virtual circuit has ceased to be the current queue, the state machine instead enqueues the cell in the current queue and henceforth changes the designation of the serving queue of this virtual circuit to make the current queue the serving queue of this virtual circuit. For example, if a cell of VC 12, whose serving queue is Q0, becomes available when Q2 is the current queue, the cell is enqueued in Q2 and VC 12's serving queue is henceforth changed from Q0 to Q2, as shown in FIG. 2.

[0024] All of this ensures that a minimum interval of 5.785ms is maintained between successive cells of the same virtual circuit output to physical layer 120, and hence ensures that a virtual circuit does not exceed its peak instantaneous cell rate of one cell per 5.785ms.

[0025] State machines 140-141 together comprise a plurality of functions, shown in FIGS. 3-6. A global initialization function of FIG. 3 initializes state machines 140-141. Upon its invocation at start-up of the system, at step 300, the function sets a global variable "emit count" 310 equal to zero, at step 302, and sets a global variable "emit Q" 312 to point to the first queue 131 of the sequence 130 of shaper queues, at step 304. Its operation then ends, at step 306.

[0026] An instance initialization function of FIG. 4 initializes an instance of enqueue state machine 140 for a particular virtual circuit. Upon its invocation at creation of a virtual circuit, at step 400, the function is passed the VPI/VCI of this virtual circuit. In response, it associates a new local variable "emit" 420 with this VPI/VCI, at step

402. "Emit" 420 is a special variable in that its least-significant bit 421 and its most-significant (eighth) bit 422 serve as flags. A set value of LSB 421 serves as a flag that indicates a first cell for this virtual circuit, and a set value of MSB 422 serves as a flag that indicates that a cell is late. After step 402, the function sets LSB 421 to a value of one, at step 404. Its operation then ends, at step 406.

[0027] An enqueueing function of FIG. 5 enqueues a cell to one of the shaper queues 131. Its execution is triggered by receipt of a cell from ATM layer 121, at step 500. Upon its invocation, it uses the VPI/VCI of the cell's ATM header to identify the "emit" variable 420 that is associated with this virtual circuit, at step 502. The function then sets a local variable "offset" 550 to equal the value of "emit" 420 minus the value of "emit count" 310, modulo 47-2 (i.e., 94), at step 504, and checks if the result indicates a first cell or a late cell for this virtual circuit, at step 506. As in the case of "emit" 420 a set LSB 421 of "offset" 550 indicates a first cell, and a set MSB 422 of "offset" 550 indicates a late cell. If neither a first nor a late cell is indicated, the function sets a local variable "Q pointer" 522 to point to the shaper queue 131 that is indexed by the value of "emit" 420 divided by two, at step 508. If either a first or a late cell is indicated at step 506, the function sets the value of "Q pointer" 522 to the value of "emit Q" 312, at step 510, and then sets the value of "emit" 420 to the value of "emit count" 310, at step 512. Following step 508 or 512, the function enqueues the cell that it received at step 500 in the shaper queue 131 pointed to by "Q pointer" 522, at step 514. It then ends its operation, at step 516, until the next cell is received.

[0028] A dequeuing function of FIG. 6 dequeues cells from one of the shaper queues 131 during each processing interval of 125us. Upon its invocation at the start of each processing interval, at step 600, the function checks whether the one of the shaper queues 131 that is pointed to by "emit Q" 312 is empty, at step 602. If it is not empty, the function dequeues a cell from that queue 131, at step 604, and sends the dequeued cell to physical layer 120, at step 606. The function then returns to step 602. If and when the function determines at step 602 that the shaper queue 131 pointed to by "emit Q" 312 is empty, its job for this operating interval is done. It increments "emit count" 310 by two, modulo 47-2 (i.e., 94), at step 608 - incrementing and counting "emit count" 310 by two ensures an emit count that is always even and thus does not mask the "first cell" flag 421 of "emit" 420. The function also increments "emit Q" 312 by one, modulo 47, to set it pointing to the next sequential one of the shaper queues 131, at step 610. The function then ends its operation, at step 612.

[0029] A second illustrative traffic shaper 121 embodying the invention is shown in FIG. 7. It is tailored to serve both narrowband and wideband CBR virtual circuits. With respect to narrowband channels, it is configured and operates as shown in FIG. 2 with the exception

that, instead of using 47 shaper queues 130, it uses N shaper queues, where N is at least equal to the number of queue processing intervals in the maximum expected cell delay (the time difference between receipt and transmittal of a cell) that traffic shaper 121 will have to introduce for a constant bit rate virtual circuit. For example, if the maximum expected cell delay is 20ms, then $N = 20\text{ms} / 125\mu\text{s} = 160$. Alternatively, N may be viewed as the multiple of the maximum expected cell delay and the peak instantaneous cell rate of the fastest one of the virtual circuits.

[0030] As was mentioned previously, each wideband circuit comprises a plurality of narrowband circuits. Hence, each wideband circuit has a plurality of corresponding serving queues—as many as there are narrowband circuits that make up this wideband circuit—evenly spaced (separated by an equal number of queues) within sequence 130 of queues 131. As shown in FIG. 7, for example, serving queues of VC 23 include Q0 and Q(N-1); serving queues of VC 20 include Q1 and Q5; serving queues of VC 22 include Q1 and Q3; serving queues of VC 25 include Q2 and Q(N-1); and serving queues of VC 25 include Q3 and Q(N-1). As in the example of FIG. 2, dequeue state machine 141 empties one of the queues 131—preferably the current queue—during each queue-processing period, and successive ones of the queues 131 are designated as the current queue during successive queue-processing periods.

[0031] When ATM layer 121 makes an ATM cell of a wideband virtual circuit available on time, i.e., while one of the serving queues of the virtual circuit is the current queue, the virtual circuit's corresponding state machine instance enqueues the cell in the current queue. For example, if a cell of VC 25 becomes available while Q2, one of VC 25's serving queues, is the current queue, the cell is enqueued in Q2, as shown in FIG. 2. If the cell is made available early, i.e., before a next serving queue of this virtual circuit becomes the current queue, the state machine instance enqueues the cell in the next serving queue of this virtual circuit. For example, if a cell of VC 24 becomes available while Q1 is the current queue, the cell is enqueued in Q2, one of VC 24's serving queues; similarly, if a cell of VC 26 becomes available while Q2 is the current queue, the cell is enqueued in Q3, one of VC 26's serving queues, as shown in FIG. 7. If the cell is made available late, i.e., after a serving queue of this virtual circuit has ceased to be the current queue, the state machine instance enqueues the cell in the current queue and the definition of each one of this virtual circuit's serving queues is shifted by the same amount so that the current queue becomes one of the virtual circuit's serving queues. For example, if a cell of VC 22, whose serving queues include Q1 and Q3, becomes available when Q2 is the current queue, the cell is enqueued in Q2 and the definition of what are VC 22's serving queues is shifted by one queue so that Q2 becomes a serving queue of VC 22 instead of Q1, Q4 becomes a serving queue of VC 22 instead of Q3, and so

on. The current queue thus becomes a substitute for the serving queue that has last ceased to be the current queue.

[0032] All of this again ensures that a minimum interval is maintained between successive cells of the same VC received by physical layer 120, and hence ensures that a virtual circuit does not exceed its peak allowed instantaneous cell rate.

[0033] State machines 140-141 of the implementation of FIG. 7 also include a plurality of functions, shown in FIGS. 3 and 8-10. Global initialization of state machines 140-141, shown in FIG. 3, does not change from the previously-described embodiment. As shown in FIG. 8, initialization of an instance of enqueue state machine 140 for a particular virtual circuit, be it a narrowband or a wideband circuit, requires the definition of an additional variable that defines the minimum allowed period of time between successive cells of that virtual circuit, i.e., it defines the spacing between serving queues of this virtual circuit in sequence 130 of shaper queues 131. For example, in FIG. 7 the "period" for a narrowband channel is N, and for a wideband channel consisting of X narrowband channels it is N/X. Upon its invocation, at step 800, the function is passed the VPI/VCI of this virtual circuit and the peak allowed instantaneous cell rate. In response, it associates new local variables "emit" 420 and "period" 820 with this VPI/VCI, at step 802. It then sets the value of "emit" 420 to one, at step 804, which serves as a flag that indicates a first cell for this virtual circuit. The function also sets the value of "period" 820, at step 806. The function then ends its operation, at step 808.

[0034] The function of FIG. 9 enqueues a cell to one of the shaper queues 131. Its execution is triggered by receipt of a cell from ATM layer 120, at step 900. Upon its invocation, it uses the VPI/VCI of the cell's ATM header to identify the "emit" variable 420 and "period" variable 820 that are associated with this virtual circuit, at step 902. The function then sets a local variable "offset" 650 to equal the value of "emit" 420 minus the value of "emit count" 310, at step 904, and checks if the result indicates a first cell or a late cell for this virtual circuit, at step 906. As in the case of "emit" 420, a set LSB 421 of "offset" 650 indicates a first cell and a set MSB 422 of "offset" 650 indicates a late cell. If neither a first nor a late cell is indicated, the function sets a local variable "Q pointer" 652 to point to the shaper queue 131 that is indexed by a value of "emit" 420 divided by two, modulo P, at step 908, where P is an integral multiple of N and $P \geq 4N$. The function then increments the value of "emit" 420 by the value of "period" 820, at step 909. If either a first or a late cell is indicated at step 906, the function sets the value of "Q pointer" 652 to the value of "emit Q" 312, at step 910, and then sets the value of "emit" 420 to the value of "emit count" 310, at step 912. Following step 909 or 912, the function enqueues the cell that it received at step 900 in the shaper queue 131 pointed to by "Q pointer" 652, at step 914. It then ends its oper-

ation, at step 916, until the next cell is received.

[0035] The function of FIG. 10 dequeues cells from one of the shaper queues 131 during each cell-processing period. Upon its invocation at the start of each period, at step 1000, the function checks whether the one of the shaper queues 131 that is pointed to by "emit Q" 312 is empty, at step 1002. If it is not empty, the function dequeues a cell from that queue 131, at step 1004, and sends the dequeued cell to physical layer 120, at step 1006. The function then returns to step 1002. If and when the function determines at step 1002 that the shaper queue 131 pointed to by "emit Q" 312 is empty, its job for this operating interval is done. It increments "emit count" 310 by two, modulo P, at step 1008, and increments "emit Q" 312 by one, modulo P, to set it pointing to the next one of shaper queues 131 in sequence 130, at step 1010. The function then ends its operation, at step 1012.

[0036] A third illustrative traffic shaper 121 embodying the invention is shown in FIG. 11. It is tailored to serve virtual circuits of any kind, be they narrowband or wideband, constant bit-rate or variable bit-rate. With respect to constant-bit-rate virtual circuits, it is configured and operates as shown in FIG. 7 and described in conjunction therewith. With respect to variable-bit-rate virtual circuits, it also is configured and operates as shown in FIG. 7 with the exception that, in addition to shaping queues 131, it makes use of an overflow queue 1100 when necessary.

[0037] A variable-bit-rate virtual circuit can produce a plurality of cells substantially at the same time. Each one of a plurality of early cells is treated as an early cell, and each is enqueued in a different one of the virtual circuit's serving queues. For example, if, during the time that Q1 is the current queue, two cells arrive for VC 34, whose serving queues include Q2 and Q4, one of the cells is enqueued in Q2 and the other cell is enqueued in Q4, as shown in FIG. 11. Each one of a plurality of cells that includes a late cell results in the late (first) cell being treated as a late cell, and the other cells are treated as early cells. For example, if, during the time that Q2 is the current queue, a late cell and an on-time cell arrive for VC 32, whose serving queues include Q1 and Q3, the late cell is enqueued in current Q2, the serving queues of VC 32 are shifted from Q1, Q3, etc. to Q2, Q4, etc., and the second cell is enqueued in Q4. There is no change in treatment of on-time cells, as only one cell can be on-time during a queue processing period.

[0038] Overflow of queues 131 occurs when so many cells arrive for a virtual circuit so quickly in succession that each serving queue of that virtual circuit presently holds one cell of that virtual circuit and the enqueueing gets ahead of the dequeuing, so that enqueueing of another cell of this virtual circuit would result either in a serving queue of this virtual circuit simultaneously storing two or more cells of this virtual circuit or in the enqueueing wrapping around the sequence of serving queues and thereby possibly causing cells to be trans-

mitted out of order. Overflow queue 1100 comes into play then, and is used to temporarily store any overflow cells of any and all virtual circuits before the dequeuing function makes room for them in queues 131. For example, as shown in FIG. 11, overflow queue 1100 may at one time hold a single overflow cell for VC 30, followed by three overflow cells for VC 35, etc., and ending in two overflow cells for VC 33. The overflow cells are periodically moved from overflow queue 1100 to queues 131 as room becomes available. To simplify operation, any cells for a virtual circuit that arrive while any cells for this virtual circuit remain in overflow queue 1100 are treated as overflow cells and are enqueued in overflow queue 1100.

[0039] State machines 140-141 of the implementation of FIG. 11 include the initialization functions shown in FIGS. 3 and 8, plus the enqueueing function shown in FIG. 12 and the dequeuing function shown in FIG. 13. Execution of the enqueueing function of FIG. 12 is triggered by the ATM layer making a cell available, at step 1200. Upon its invocation, it uses the VPI/VC1 of the cell's ATM header to identify the "emit" variable 420 and "period" variable 820 that are associated with this virtual circuit, at step 1202. The function then sets a local variable "offset" 650 to equal the value of "emit" 420 minus the value of "emit count" 310, at step 904, and checks if the result is greater than N (the number of shaping queues 131 in sequence 130), at step 1205. Value of "offset" 650 greater than N indicates an overflow condition, and so the function merely enqueues the received cell on overflow queue 1100, at step 1218, and then ends its operation, at step 1220. If an overflow condition is not detected at step 1205, the function proceeds to steps 1206 et seq. that duplicate the steps 906 et seq. of FIG. 9, with the exception that the modulo operation in steps 1209 and 1212 is modulo R and not modulo P. R is an integral multiple of N at least equal to P, and is also equal to or greater than the number of cell processing periods in the maximum expected cell delay that traffic shaper 121 will have to introduce for any virtual circuit. For example, if N is 160 and the maximum expected cell delay for any virtual circuit is 70ms, then $R = \lceil 70\text{ms} / 125\mu\text{s} \rceil_N = 4N = 640$.

[0040] The function of FIG. 13 dequeues cells from overflow queue 1100 once during each full cycle of operating intervals and dequeues calls from one of the shaper queues during each operating interval. Upon its invocation at the start of each interval, at step 1300, the function checks whether the value of "emit count" 310 is zero, at step 1301. If so, it means that a new cycle of operating intervals is starting, and so the function proceeds to do overflow queue 1100 processing. It checks whether overflow queue 1100 is empty, at step 314. If it is not empty, the function dequeues a cell from overflow queue 1100, at step 1316, and sends it to the enqueueing function of FIG. 12, at step 1318. The enqueueing function of FIG. 12 treats the cell as if it had just been received from ATM layer 122. The dequeuing function

then returns to step 314. If and when the dequeuing function finds at step 314 that overflow queue 1100 is empty, or if it finds at step 1301 that "emit count" 310 is not zero, it proceeds to steps 1302 et seq. which duplicate the steps 1002 et seq. of FIG. 10, with the exception that the modulo operation in step 1308 is modulo R and not modulo P.

[0041] Of course, various changes and modifications to the illustrative embodiments described above will be apparent to those skilled in the art. For example, the overflow queue may take the form of a hierarchy of a plurality of overflow queues, with each level in the hierarchy corresponding to a different number of processing cycles (processing cycle=processing interval-number of shaper queues) and for storing cells arriving the corresponding number of processing cycles early. Such changes and modifications can be made within the scope of the invention and without diminishing its attendant advantages. It is therefore intended that such changes and modifications be covered by the following claims except insofar as limited by the prior art.

Claims

1. A traffic shaper (121) for a plurality of streams (VC) of traffic, comprising:

a first plurality of queues (131);
 dequeuing means (141) cooperative with the first plurality of queues for transmitting contents of a sequence of the queues; and
 enqueueing means (140) cooperative with the first plurality of queues and responsive to receipt of traffic from a traffic stream for storing (154) the received traffic in the serving (508) queue of the traffic stream,

CHARACTERISED IN THAT

the first plurality of queues comprises a sequence (130) of the queues shared by the plurality of traffic streams, each of the traffic streams having at least one of the queues designated (420, 550) as its serving queue;
 the dequeuing means cyclically transmit (606) the contents of a different sequential (610, 312) one of the queues during each sequential period of a said first plurality of periods, wherein during each period a queue having a fixed offset in the sequence from the queue whose contents are being transmitted forms a current queue for that period; and
 the enqueueing means are responsive to receipt (500) of traffic from any said traffic stream during a period before a serving queue of that traffic stream becomes the current queue for said period, for storing (514) the received traffic in

the serving (508) queue of the traffic stream, are further responsive to receipt of traffic from any said traffic stream during a period when a serving queue of that traffic stream is the current queue for said period, for storing (514) the receiving traffic in the current (508) queue for said period, and are further responsive to receipt of traffic from any said traffic stream during a period after a serving queue of that traffic stream ceased being the current queue for said period, for storing (514) the received traffic in the current (510) queue for said period and changing (512) the designation of the at least one serving queue of that traffic stream to designate the current queue of said period as a serving queue of that traffic stream.

2. A traffic shaper as claimed in claim 1 wherein:

the at least one serving queue of a traffic stream comprises a plurality of serving queues spaced (820) evenly in the sequence of queues, and the enqueueing means change the designation of the at least one serving queue of said traffic stream by shifting (912) the designation from each of the serving queues a same distance in the sequence to other queues that include the current queue.

3. A traffic shaper as claimed in claim 1 wherein:

the first plurality of queues comprises at least as many queues as there are said periods in a maximum time between receipt of traffic of any stream by the enqueueing means and transmission of said traffic of said stream by the dequeuing means; and
 each period comprises a time interval at most equal to an inverse of a peak instantaneous rate of receipt of the traffic of a fastest one of the traffic streams.

4. A traffic shaper as claimed in claim 1 wherein: the fixed offset is zero.

5. A traffic shaper as claimed in claim 1 wherein:

at least one of the traffic streams is a variable-bit-rate traffic stream of one of a plurality of different sizes;
 the one traffic stream has a non-zero number of queues proportional to its size designated (420, 820) as its at least one serving queue, and the serving queues of the one traffic stream are equally spaced within the sequence of queues; and
 the traffic shaper further comprises at least one overflow queue (1100) for storing

(1218) any traffic of any stream received when operation of the enqueueing means would get ahead (1205) of operation of the dequeuing means, and

means for causing (1316, 1318) the enqueueing means to receive traffic stored in the overflow queue when operation (1004) of the dequeuing means has made room for storing the traffic stored in the overflow queue available in the sequence of the queues.

6. A method of shaping traffic of a plurality of streams of traffic, comprising:

in response to receipt of traffic from a traffic stream, storing the received traffic in a serving queue of that traffic stream among a first plurality of queues; and
transmitting contents of a sequence of the queues,

CHARACTERISED BY

for each one of the plurality of traffic streams, designating (420, 550) at least one of the sequence (130) of the first plurality of queues (131) shared by the plurality of traffic streams as a serving queue of said one traffic stream; wherein

the transmitting comprises cyclically transmitting (606) contents of the sequence of the queues by transmitting the contents of a different sequential (610, 312) one of the queues during each sequential period of a said first plurality of periods, wherein during each period a queue having a fixed offset in the sequence from the queue whose contents are being transmitted forms a current queue for that period; and

the storing comprises in response to receipt (500) of traffic from any said traffic stream during a period before a serving queue of that traffic stream becomes the current queue for said period, storing (514) the received traffic in the serving (508) queue of that traffic stream,

in response to receipt of traffic from any said traffic stream during a period when a serving queue of that traffic stream is the current queue for said period, storing (514) the received traffic in the current (508) queue, and

in response to receipt of traffic from any said traffic stream during a period after a serving queue of that traffic stream ceased being the current queue for said period, storing (514) the received traffic in the current (510) queue for that period, and changing (512) the designation of the at least one serving queue of that traffic

stream to designate the current queue for said period as a serving queue of that traffic stream.

7. A method as claimed in claim 6 wherein:

designating comprises for at least one traffic stream, designating a plurality of the queues spaced (820) evenly in the sequence of queues as the serving queues of said one traffic stream; and
changing the designation comprises shifting (912) the designation from each of the serving queues of said one traffic stream a same distance in the sequence to other queues that include the current queue.

8. A method as claimed in claim 6 wherein:

the first plurality of queues comprises at least as many queues as there are said periods in a maximum time between receipt of traffic of any stream by the enqueueing means and transmission of said traffic of said stream by the dequeuing means; and
each period comprises a time interval at most equal to an inverse of a peak instantaneous rate of receipt of the traffic of a fastest one of the traffic streams.

9. A method as claimed in claim 6 wherein:
the fixed offset is zero.

10. A method as claimed in claim 6 wherein:

at least one of the traffic streams is a variable-bit-rate traffic stream of one of a plurality of different sizes; and
the one traffic stream has a non-zero number of queues proportional to its size designated (420, 820) as its at least one serving queue, and the serving queues of the one traffic stream are equally spaced within the sequence of queues; and
each step of storing comprises storing (1218) the received traffic in an overflow queue (1100) instead of a queue of the first plurality of queues if storing the received traffic in the queue of the first plurality of queues would get ahead (1205) of the cyclically transmitting, and
in response to room for storing the traffic stored in the overflow queue having been made available in the first plurality of queues by said cyclically transmitting (1004) contents of the sequence of queues, storing (1316, 1318) the traffic stored in the overflow queue in the first plurality of queues.

Patentansprüche

1. Verkehrsformungsvorrichtung (121) für eine Vielzahl von Verkehrsströmen (VC), aufw isend:

eine erste Vielzahl von Warteschlangen (131);
eine mit der ersten Vielzahl von Warteschlangen zusammenarbeitende Warteschlangen-Auflösungseinrichtung zum Übertragen von Inhalten einer Folge von den Warteschlangen; und
eine mit der ersten Vielzahl von Warteschlangen zusammenarbeitende und auf den Empfang von Verkehr aus einem Verkehrsstrom ansprechende Warteschlangen-Aufbaueinrichtung (140) zum Speichern (154) des empfangenen Verkehrs in der Abwicklungswarteschlange (508) des Verkehrsstroms,

dadurch gekennzeichnet, daß

die erste Vielzahl von Warteschlangen ein Folge (130) von den von der Vielzahl der Verkehrsströme gemeinsam genutzten Warteschlangen aufweist, wobei jeder Verkehrsstrom mindestens eine von den bezeichneten Warteschlangen (420, 550) als seine Abwicklungswarteschlange besitzt;
die Warteschlangen-Auflösungseinrichtung zyklisch den Inhalt einer anderen folgenden Warteschlange (610, 312) während jeder folgenden Periode der ersten Vielzahl von Perioden überträgt (606), wobei während jeder Periode eine Warteschlange mit einer festen Verschiebung in der Folge von der Warteschlange aus, deren Inhalte übertragen werden, eine aktuelle Warteschlange für diese Periode bildet; und
die Warteschlangen-Aufbaueinrichtung auf den Empfang (500) von Verkehr von jedem Verkehrsstrom während einer Periode anspricht, bevor eine Abwicklungswarteschlange dieses Verkehrsstroms die aktuelle Warteschlange für diese Periode wird, um den empfangenen Verkehr in der Abwicklungswarteschlange (508) des Verkehrsstroms zu speichern (514), und ferner auf den Empfang von Verkehr von jedem Verkehrsstrom während einer Periode anspricht, wenn eine Abwicklungswarteschlange dieses Verkehrsstroms die aktuelle Warteschlange für diese Periode ist, um den empfangenen Verkehr in der aktuellen Warteschlange (508) für diese Periode zu speichern (514), und ferner auf den Empfang von Verkehr von jedem Verkehrsstrom während einer Periode reagiert, nachdem eine Abwicklungswarteschlange dieses Verkehrsstroms aufhörte eine Abwicklungswarteschlange für diese Periode zu sein, um den empfangenen Verkehr in der aktuellen

Warteschlange (510) für diese Periode zu speichern (514) und die Bezeichnung der mindestens einen Abwicklungswarteschlange dieses Verkehrsstroms zu ändern (512), um die aktuelle Warteschlange dieser Periode als eine Abwicklungswarteschlange dieses Verkehrsstroms zu bezeichnen.

2. Verkehrsformungsvorrichtung nach Anspruch 1, wobei:

die mindestens eine Abwicklungswarteschlange eines Verkehrsstroms eine Vielzahl von Abwicklungswarteschlangen aufweist, die in der Folge von Warteschlangen gleichmäßig beabstandet sind (820) sind, und
die Warteschlangen-Aufbaueinrichtung die Bezeichnung der mindestens einen Abwicklungswarteschlange des Verkehrsstroms durch Verschiebung (912) der Bezeichnung von jeder Abwicklungswarteschlange um einen gleichen Abstand in der Folge auf andere Warteschlangen, welche die aktuelle Warteschlange enthalten, verändert.

3. Verkehrsformungsvorrichtung nach Anspruch 1, wobei:

die erste Vielzahl von Warteschlangen mindestens so viele Warteschlangen umfaßt, wie die Perioden in einer maximalen Zeit zwischen dem Empfang von Verkehr von jedem Strom durch die Warteschlangen-Aufbaueinrichtung und der Übertragung des Verkehrs des Stroms durch die Warteschlangen-Auflösungseinrichtung vorhanden sind; und
jede Periode ein Zeitintervall nahezu gleich einer Inversen der Spitzenmomentanrate eines Empfangs des Verkehrs eines schnellsten von den Verkehrsströmen aufweist.

4. Verkehrsformungsvorrichtung nach Anspruch 1, wobei:

die feste Verschiebung Null ist;

5. Verkehrsformungsvorrichtung nach Anspruch 1, wobei:

mindestens einer von den Verkehrsströmen ein Verkehrsstrom mit variabler Bitrate mit einer von einer Vielzahl unterschiedlichen Größen ist;
der eine Verkehrsstrom eine Anzahl von Warteschlangen ungleich Null proportional zu seiner angegebenen Größe (420, 480) als seine mindestens eine Abwicklungswarteschlange besitzt, und die Abwicklungswarteschlangen des einen Verkehrsstroms innerhalb der Fol-

gen der Warteschlangen gleich beabstandet sind; und
die Verkehrsformungsvorrichtung ferner umfaßt:

mindestens eine Überlaufwarteschlange (1100) zum Speichern (1218) jedes Verkehrs jedes empfangenen Stroms, wenn der Betrieb der Warteschlangen-Aufbaueinrichtung vor dem Betrieb der Warteschlangen-Auflösungseinrichtung stattfindet (1205); und
eine Einrichtung, um die Warteschlangen-Aufbaueinrichtung zu veranlassen (1316, 1318), in der Überlaufwarteschlange gespeicherten Verkehr zu empfangen, wenn die Operation (1004) der Warteschlangen-Auflösungseinrichtung Platz für die Speicherung des Verkehrs geschaffen hat, der in der in der Folge der Warteschlangen verfügbaren Überlaufwarteschlange gespeichert ist.

6. Verfahren zur Speicherung von einer Vielzahl von Verkehrsströmen, welches die Schritte aufweist:

als Antwort auf den Empfang eines Verkehrs aus einem Verkehrsstrom Speichern des empfangenen Verkehrs in einer Abwicklungswarteschlange dieses Verkehrsstroms unter einer ersten Vielzahl von Warteschlangen; und
Übertragen der Inhalte einer Folge von Warteschlangen,

dadurch gekennzeichnet, daß

für jeden einzelnen von der Vielzahl von Verkehrsströmen mindestens eine von den Folgen (130) der ersten Vielzahl von Warteschlangen (131), welche von den Vielzahl der Verkehrsströme gemeinsam genutzt werden, als eine Abwicklungswarteschlange des eines Verkehrsstroms bezeichnet wird (420, 550); wobei das Übertragen umfaßt:

ein zyklisches Übertragen (606) der Inhalte der Folge von den Warteschlangen, indem die Inhalte eine anderen folgenden Warteschlange (610, 312) während jeder folgenden Periode der ersten Vielzahl von Perioden übertragen werden, wobei während jeder Periode eine Warteschlange mit einer festgelegten Verschiebung in der Folge von der Warteschlange aus, deren Inhalte gerade übertragen werden, eine aktuelle Warteschlange für dieses Periode bildet; und
das Speichern umfaßt:

unter Ansprechen auf den Empfang von Verkehr von jedem Verkehrsstrom während einer Periode, bevor eine Abwicklungswarteschlange dieses Verkehrsstroms die aktuelle Warteschlange für diese Periode wird, das Speichern (514) des empfangenen Verkehrs in der Abwicklungswarteschlange (508) dieses Verkehrsstroms,
unter Ansprechen auf den Empfang von Verkehr von jedem Verkehrsstrom während einer Periode, wenn eine Abwicklungswarteschlange dieses Verkehrsstroms die aktuelle Warteschlange für diese Periode ist, das Speichern (514) des empfangenen Verkehrs in der aktuellen Warteschlange (508),
und
unter Ansprechen auf den Empfang von Verkehr von jedem Verkehrsstrom während einer Periode, nachdem eine Abwicklungswarteschlange dieses Verkehrsstroms aufhörte eine Abwicklungswarteschlange für diese Periode zu sein, das Speichern (514) des empfangenen Verkehrs in der aktuellen Warteschlange (510) für diese Periode und das Verändern (512) der Bezeichnung der mindestens einen Abwicklungswarteschlange dieses Verkehrsstroms, um die aktuelle Warteschlange dieser Periode als eine Abwicklungswarteschlange dieses Verkehrsstroms zu bezeichnen.

7. Verfahren nach Anspruch 6, wobei das Bezeichnen umfaßt:

für mindestens einen Verkehrsstrom das Bezeichnen einer Vielzahl von Warteschlangen, die gleichmäßig in der Folge von Warteschlangen beabstandet (820) sind, als die Abwicklungswarteschlangen des einen Verkehrsstroms; und

das Ändern der Bezeichnung umfaßt:

das Verschieben (912) der Bezeichnung von jeder Abwicklungswarteschlange des einen Verkehrsstroms um den gleichen Abstand in der Folge auf andere Warteschlangen, welche die aktuelle Warteschlange enthalten.

8. Verfahren nach Anspruch 6, wobei

die erste Vielzahl von Warteschlangen mindestens so viele Warteschlangen umfaßt, wie Perioden in einer maximalen Zeit zwischen dem Empfang von Verkehr jedes Stroms durch die Warteschlangen-Aufbaueinrichtung und der

Übertragung des Verkehrs des Stroms durch die Warteschlangen-Auflösungseinrichtung vorhanden sind; und
jede Periode ein Zeitintervall nahezu gleich einer Inversen der Spitzenmomentanrate eines Empfangs des Verkehrs eines schnellsten von den Verkehrsströmen umfaßt.

9. Verfahren nach Anspruch 6, wobei die feste Verschiebung Null ist;

10. Verfahren nach Anspruch 6, wobei

mindestens einer von den Verkehrsströmen ein Verkehrsstrom mit variabler Bitrate mit einer von einer Vielzahl unterschiedlichen Größen ist;
der eine Verkehrsstrom eine Anzahl von Warteschlangen ungleich Null proportional zu seiner angegebenen Größe (420, 480) als seine mindestens eine Abwicklungswarteschlange besitzt, und die Abwicklungswarteschlangen des einen Verkehrsstroms innerhalb der Folgen der Warteschlangen gleich beabstandet sind; und
jeder Speicherungsschritt umfaßt:

das Speichern (1218) des empfangenen Verkehrs in einer Überlaufwarteschlange (1100) anstelle einer Warteschlange der ersten Vielzahl von Warteschlangen, falls die Speicherung des empfangenen Verkehrs in der Warteschlange der ersten Vielzahl von Warteschlangen vor der zyklischen Übertragung stattfindet (1205), und als Antwort auf Platz zum Speichern des in der Überlaufwarteschlange gespeicherten Verkehrs, der in der ersten Vielzahl von Warteschlangen durch das zyklische Übertragen (1004) von Inhalten der Folge von Warteschlangen verfügbar gemacht wurde, das Speichern (1316, 1318) des in der Überlaufwarteschlange gespeicherten Verkehrs in der ersten Vielzahl von Warteschlangen.

Revendications

1. Dispositif de mise en forme du trafic (121) destiné à une pluralité de courants (VC) de trafic, comprenant :

une première pluralité de files d'attente (131) ; des moyens de sortie de file d'attente (141) coopérant avec la première pluralité de files d'attente pour transmettre des contenus d'une séquence des files d'attente ; et

des moyens de mise en file d'attente (140) coopérant avec la première pluralité de files d'attente et sensibles à la réception de trafic d'un courant de trafic pour stocker (154) le trafic reçu dans la file d'attente de desserte (508) du courant de trafic,

caractérisé en ce que :

la première pluralité de files d'attente comprend une séquence (130) des files d'attente partagées par la pluralité de courants de trafic, chacun des courants de trafic ayant au moins l'une des files d'attente désignée (420, 550) en tant que sa file d'attente de desserte ;
les moyens de sortie de file d'attente transmettant de façon cyclique (606) les contenus de l'une différente et séquentielle (610, 312) des files d'attente pendant chaque période séquentielle de la première pluralité de périodes, dans lequel, pendant chaque période, une file d'attente ayant un décalage fixe dans la séquence par rapport à la file d'attente dont les contenus sont en cours de transmission, forme une file d'attente courante pour cette période ; et
les moyens de mise en file d'attente sont sensibles à la réception (500) de trafic de l'un quelconque desdits courants de trafic pendant une période précédant le moment où une file d'attente de desserte de ce courant de trafic devient la file d'attente courante pour cette période, pour stocker (514) le trafic reçu dans la file d'attente de desserte (508) du courant de trafic, sont en outre sensibles à la réception de trafic de l'un quelconque des courants de trafic pendant une période lors de laquelle une file d'attente de desserte de ce courant de trafic est la file d'attente courante pour la période en question, pour stocker (514) le trafic de réception dans la file d'attente courante (508) pour la période en question, et sont en outre sensibles à la réception de trafic de l'un quelconque des courants de trafic pendant une période faisant suite au moment où une file d'attente de desserte de ce courant de trafic a cessé d'être la file d'attente courante pour la période en question, pour stocker (514) le trafic reçu dans la file d'attente courante (510) pour cette période et modifier (512) la désignation de cette au moins une file d'attente de desserte de ce courant de trafic pour désigner la file d'attente courante de la période en question comme étant la file d'attente de desserte de ce courant de trafic.

2. Dispositif de mise en forme de trafic suivant la revendication 1, dans lequel :

cette au moins une file d'attente de desserte

- d'un courant de trafic comprend une pluralité de files d'attente de desserte espacées (520) régulièrement dans la séquence de files d'attente, et
les moyens de mise en file d'attente modifient la désignation de cette au moins une file d'attente de desserte du courant de trafic en décalant (912) la désignation de chacune des files d'attente de desserte d'une même distance dans la séquence vers d'autres files d'attente qui incluent la file d'attente courante.
3. Dispositif de mise en forme de trafic suivant la revendication 1, dans lequel :
- la première pluralité de files d'attente comprend au moins autant de files d'attente qu'il y a de périodes en question en un temps maximum entre la réception de trafic de l'un quelconque des courants par les moyens de mise en file d'attente et la transmission du trafic de ce courant par les moyens de sortie de file d'attente ; et
chaque période comprend un intervalle de temps au moins égal à un inverse du débit instantané crête de réception du trafic du plus rapide des courants de trafic.
4. Dispositif de mise en forme de trafic suivant la revendication 1, dans lequel :
le décalage fixe est nul.
5. Dispositif de mise en forme de trafic suivant la revendication 1, dans lequel :
- au moins l'un des courants de trafic est un courant de trafic à débit binaire variable ayant l'une d'une pluralité de tailles différentes ;
le courant de trafic en question a un nombre non nul de files d'attente proportionnel à sa taille désignée (420, 820) en tant que sa au moins une file d'attente de desserte, et les files d'attente de desserte du courant de trafic en question sont régulièrement espacées dans la séquence de files d'attente ; et
le dispositif de mise en forme de trafic comprend en outre
au moins une file d'attente de débordement (1100) pour stocker (1218) tout trafic de tout courant reçu lorsque le fonctionnement des moyens de mise en file d'attente prend de l'avance (1205) sur le fonctionnement des moyens de sortie de file d'attente, et
des moyens pour faire en sorte (1316, 1318) que les moyens de mise en file d'attente reçoivent le trafic stocké dans la file d'attente de débordement lorsque le fonctionnement (1004) des moyens de sortie de file d'attente a libéré de la place pour le stockage du trafic stocké dans la file d'attente de débordement disponible dans la séquence de files d'attente.
6. Procédé de mise en forme de trafic d'une pluralité de courants de trafic, comprenant les étapes qui consistent à :
- en réponse à la réception de trafic d'un courant de trafic, stocker le trafic reçu dans une file d'attente de desserte de ce courant de trafic parmi une première pluralité de files d'attente ; et
transmettre des contenus d'une séquence de files d'attente,
- caractérisé par les étapes qui consistent à :
- pour chacun de la pluralité de courants de trafic, désigner (420, 550) au moins une file de la séquence (130) de la première pluralité de files d'attente (131) partagées par la pluralité de courants de trafic en tant que file d'attente de desserte du courant de trafic en question ; dans lequel
la transmission comprend
la transmission cyclique (606) de contenus de la séquence des files d'attente en transmettant les contenus d'une file différente et séquentielle (610, 312) des files d'attente pendant chaque période séquentielle de la première pluralité de périodes lors de laquelle une file d'attente ayant un décalage fixe dans la séquence par rapport à la file d'attente dont les contenus sont en cours de transmission, forme une file d'attente courante pour cette période ; et
le stockage comprend,
en réponse à la réception (500) de trafic de l'un quelconque des courants de trafic pendant une période précédant le moment où une file d'attente de desserte de ce courant de trafic devient la file d'attente courante pour cette période, le stockage (514) du trafic reçu dans la file d'attente de desserte (508) de ce courant de trafic,
en réponse à la réception de trafic de l'un quelconque des courants de trafic pendant une période lors de laquelle une file d'attente de desserte de ce courant de trafic est la file d'attente courante pour cette période, le stockage (514) du trafic reçu dans la file d'attente courante (508), et
en réponse à la réception de trafic de l'un quelconque des courants de trafic pendant une période faisant suite au moment où une file d'attente de desserte de ce courant de trafic a cessé d'être la file d'attente courante pour cette période, le stockage (514) du trafic reçu dans la file d'attente (510) courante pour cette période,

et la modification (512) de la désignation de cette au moins une file d'attente de desserte de ce courant de trafic pour désigner la file d'attente courante pour la période en question comme étant une file d'attente de desserte de ce courant de trafic.

5

7. Procédé suivant la revendication 6, dans lequel :

la désignation comprend,
pour au moins un courant de trafic, la désignation d'une pluralité de files d'attente espacées (820) régulièrement dans la séquence de files d'attente en tant que files d'attente de desserte du courant de trafic en question ; et
la modification de la désignation comprend le décalage (912) de la désignation de chacune des files d'attente de desserte du courant de trafic en question, d'une même distance dans la séquence vers d'autres files d'attente qui contiennent la file d'attente courante.

10

15

20

8. Procédé suivant la revendication 6, dans lequel :

la première pluralité de files d'attente comprend au moins autant de files d'attente qu'il y a de périodes en question en un temps maximum entre la réception de trafic d'un courant quelconque par les moyens de mise en file d'attente et la transmission de ce trafic de ce courant par les moyens de sortie de file d'attente ; et
chaque période comprend un intervalle de temps au plus égal à un inverse du débit instantané crête de réception du trafic d'un plus rapide des courants de trafic.

25

30

35

9. Procédé suivant la revendication 6, dans lequel :
le décalage fixe est nul.

10. Procédé suivant la revendication 6, dans lequel :

40

au moins l'un des courants de trafic est un courant de trafic à débit binaire variable ayant l'une d'une pluralité de tailles différentes ; et
le courant de trafic en question a un nombre non nul de files d'attente proportionnel à sa taille désignée (420, 820) comme étant sa au moins une file d'attente de desserte, et les files d'attente de desserte du courant de trafic en question sont régulièrement espacées dans la séquence de files d'attente ; et
chaque étape de stockage comprend le stockage (1218) du trafic reçu dans une file d'attente de débordement (1100) au lieu d'une file d'attente de la première pluralité de files d'attente si le stockage du trafic reçu dans la file d'attente de la première pluralité de files d'attente prend de l'avance (1205) sur la trans-

45

50

55

mission cyclique, et
en réponse au fait que de la place pour le stockage du trafic stocké dans la file d'attente de débordement a été rendue disponible dans la première pluralité de files d'attente par cette transmission cyclique (1004) des contenus de la séquence de files d'attente, le stockage (1316, 1318) du trafic stocké dans la file d'attente de débordement dans la première pluralité de files d'attente.




FIG. 1




FIG. 2

FIG. 3




FIG. 4




FIG. 5





FIG. 6






FIG. 7

FIG. 8

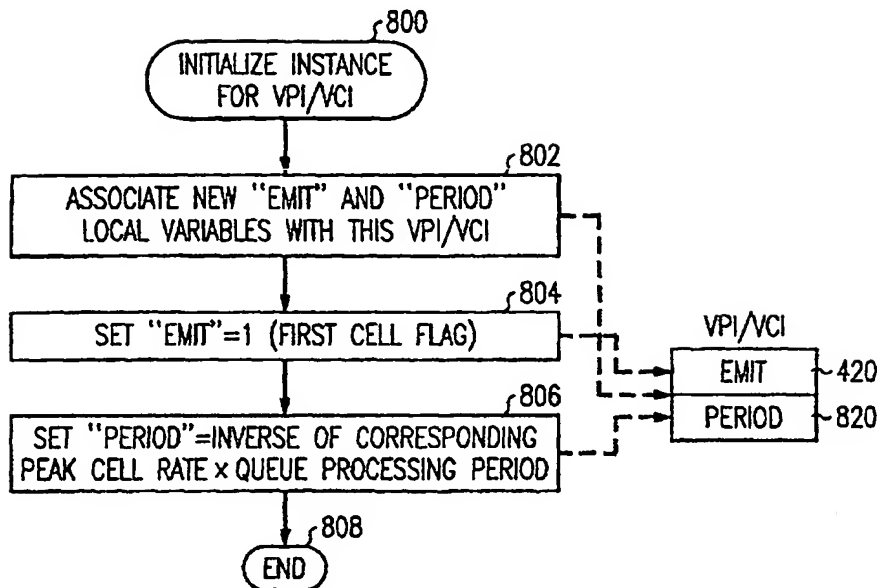


FIG. 10


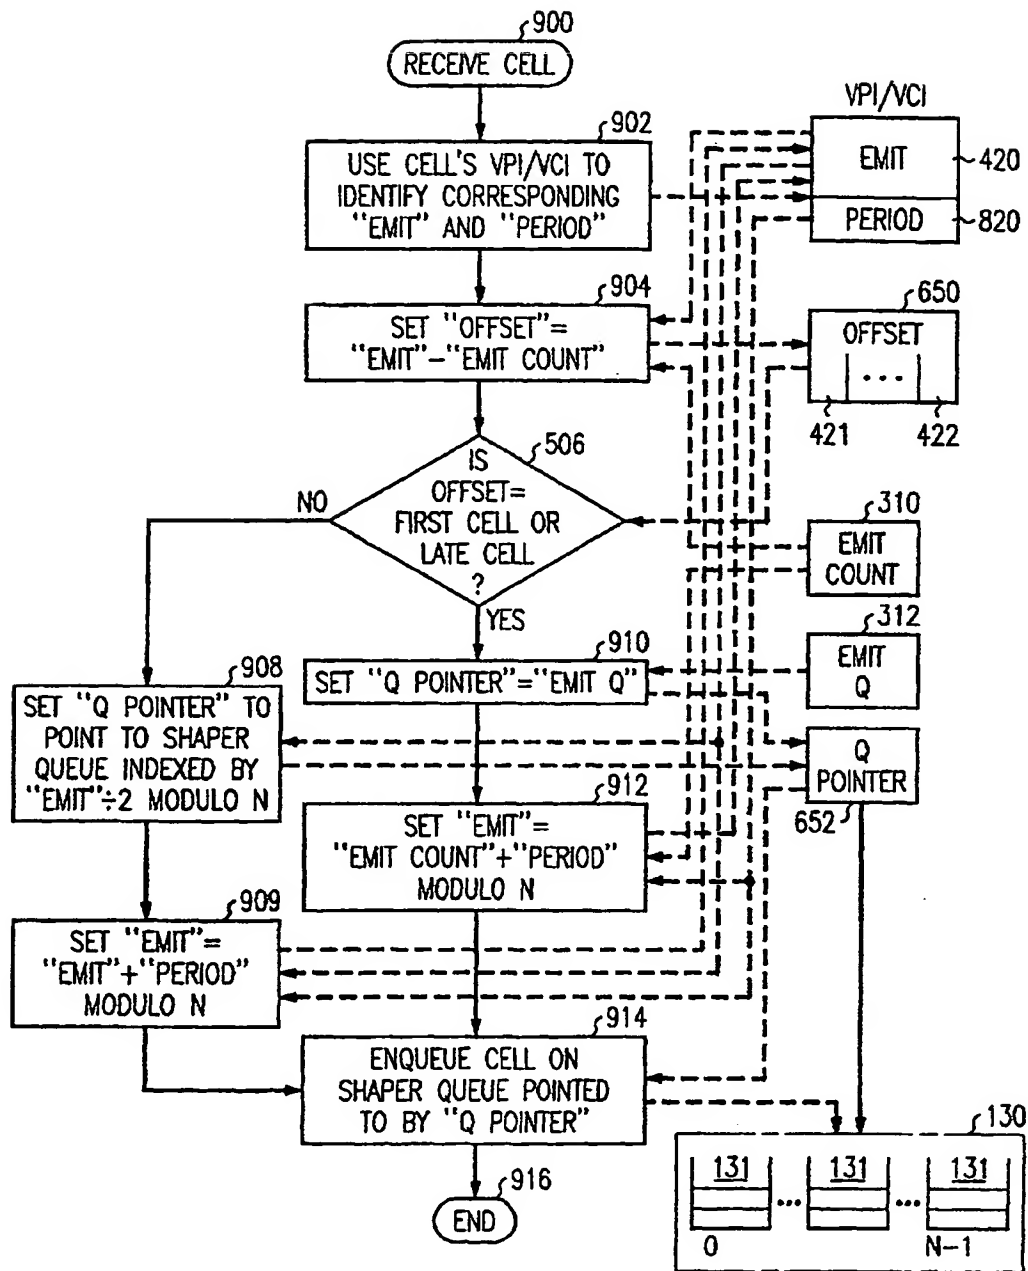


FIG. 9






FIG. 11

FIG. 12




FIG. 13

